

facebook

OpenBMC

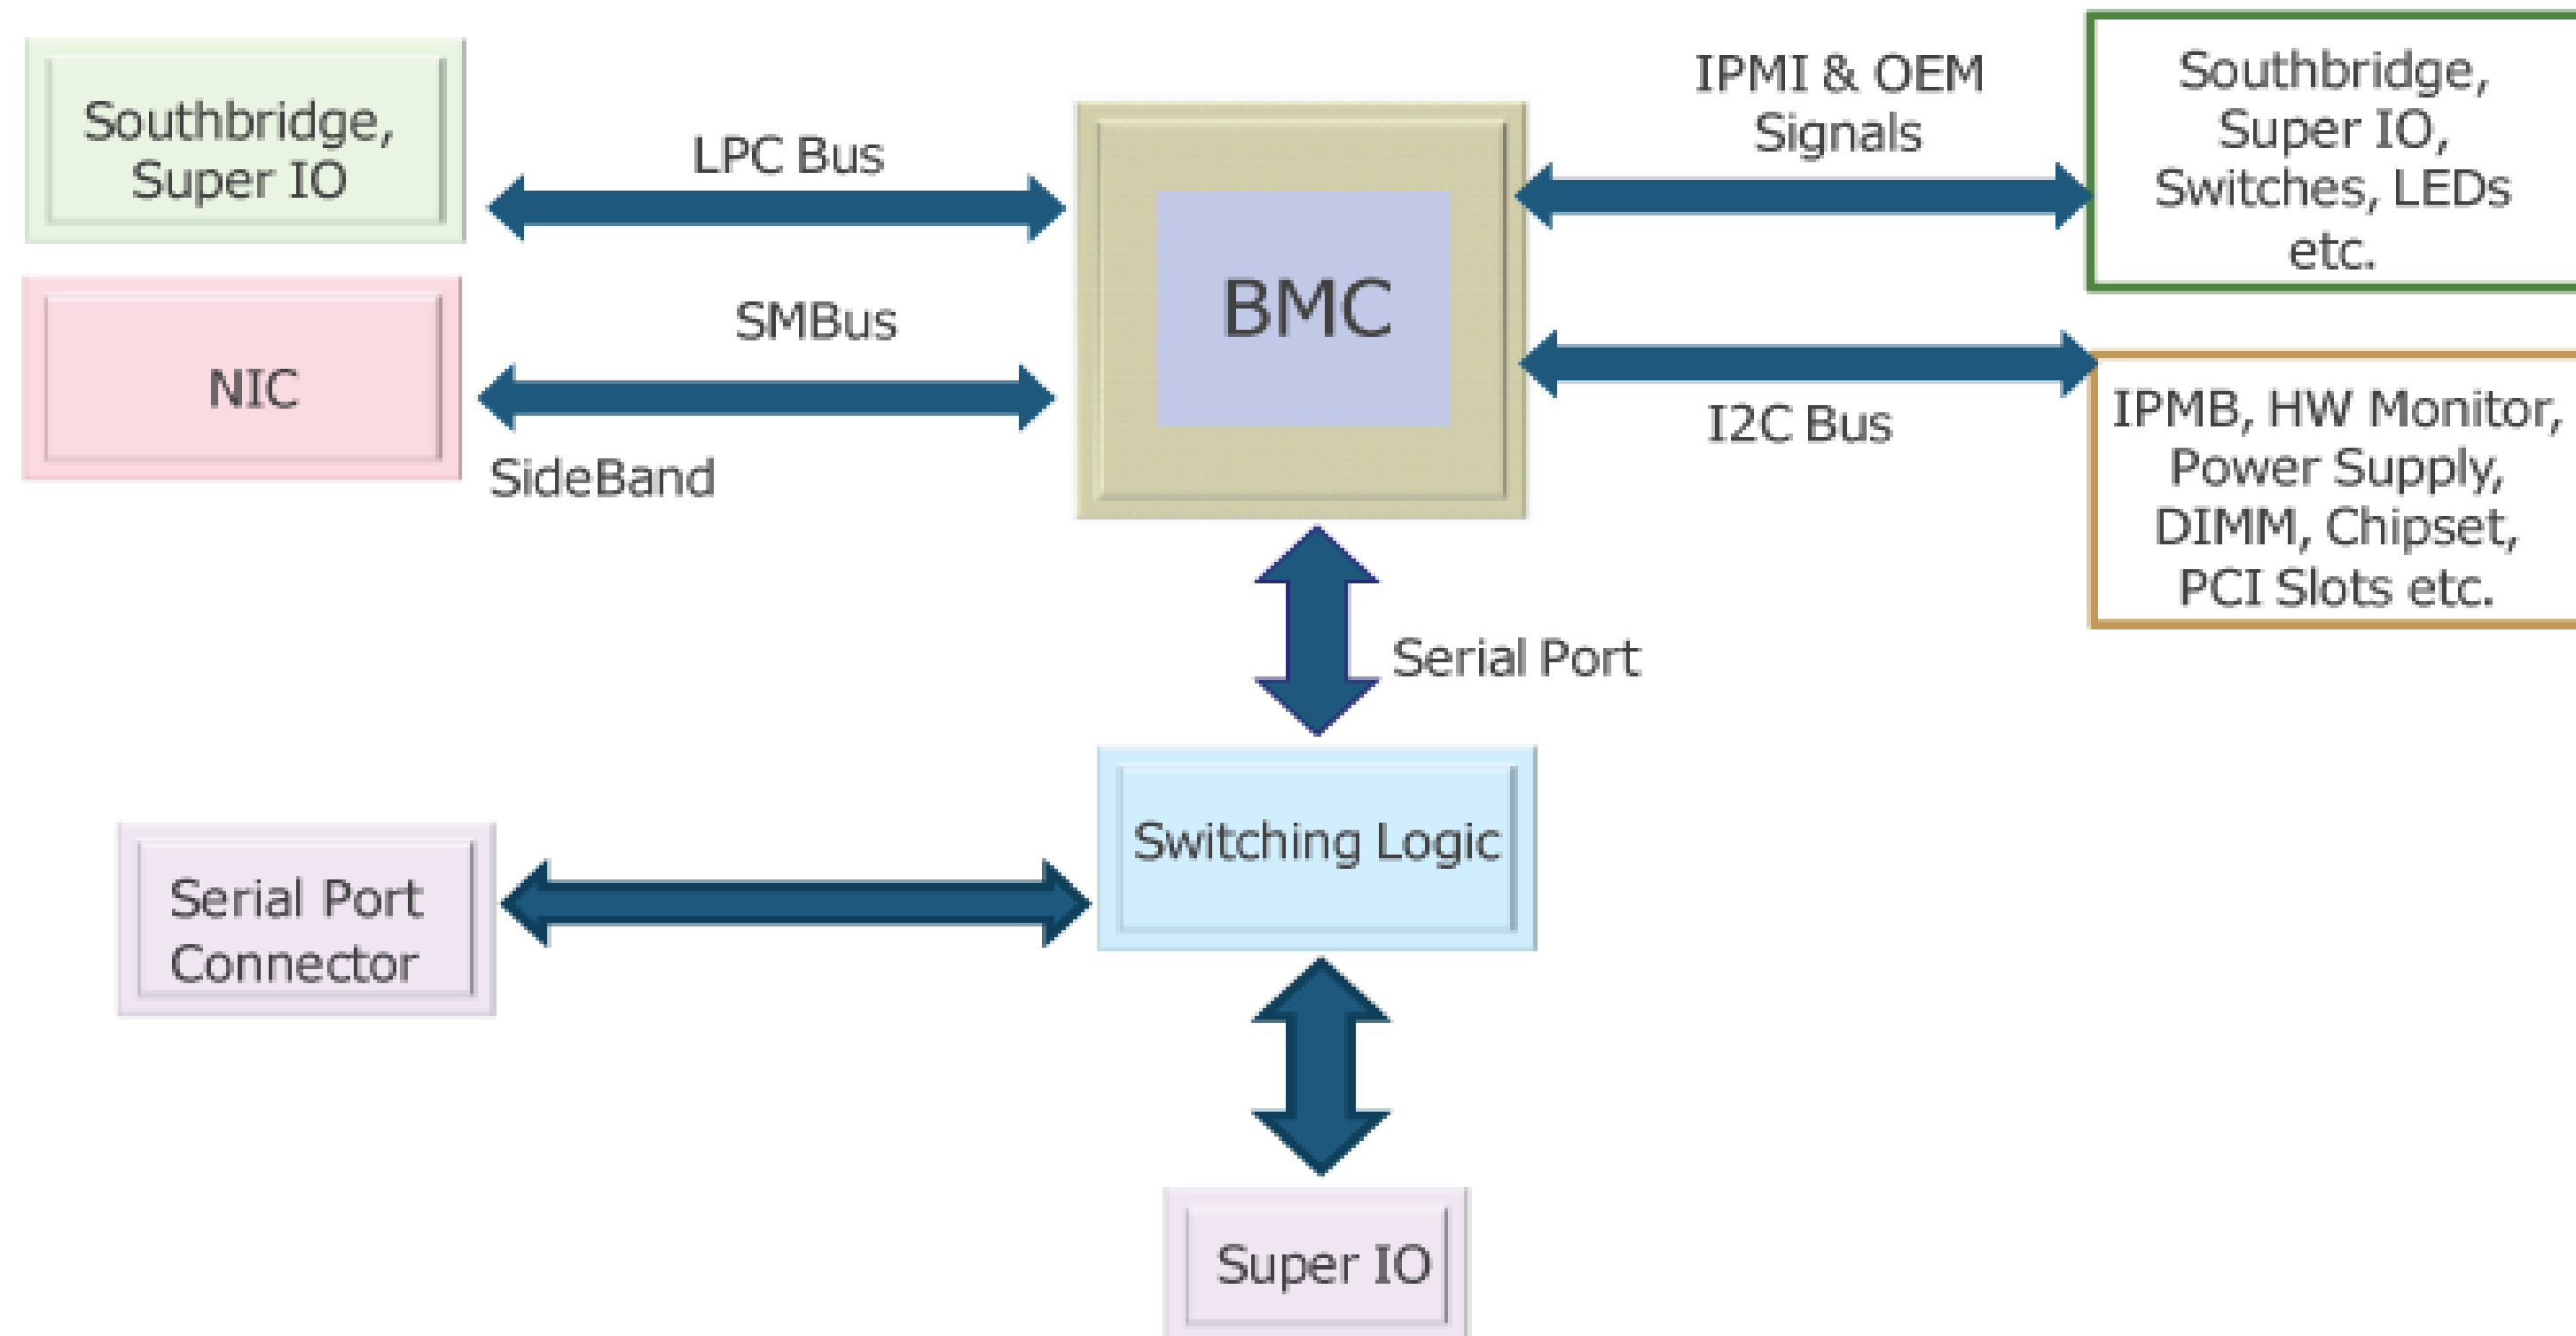
Tian Fang

Software Engineer - Facebook

Baseboard Management Controller

Intelligent Platform Management Interface (IPMI)

IPMI Block Diagram



Feature Comparisons





	AST2500	AST2400	AST2300
System Bus Interface	PCIe 1x, gen 2	PCIe 1x	PCIe 1x
PCIe Host	Yes	no	no
VGA/2D Controller	Yes	Yes	Yes
BMC	Yes	Yes	Yes
Storage Re-Direction	Yes	Yes	Yes
KVM Re-Direction	Yes	Yes	Yes
ARM926EJ CPU Speed	800MHz, ARM 11	400MHz	400MHz
32bits ColdFire V1 CPU	200MHz	200MHz	Disabled
SDRAM Type	DDR4/DDR3 LV	DDR3/DDR2	DDR3/DDR2
SDRAM Speed	DDR4 1600	DDR3 800	DDR3 800
SDRAM Bus Width	16-Bit	16-Bit	16-Bit
Max. SDRAM Capacity	1GB	512MB	512MB
SDRAM ECC Capability	Yes (1/8 Size)	Yes (1/8 Size)	Yes (1/8 Size)
Flash Memory Type	SPI x3	SPI x2	SPI x2
Last Frame Capture	Yes	Yes	No

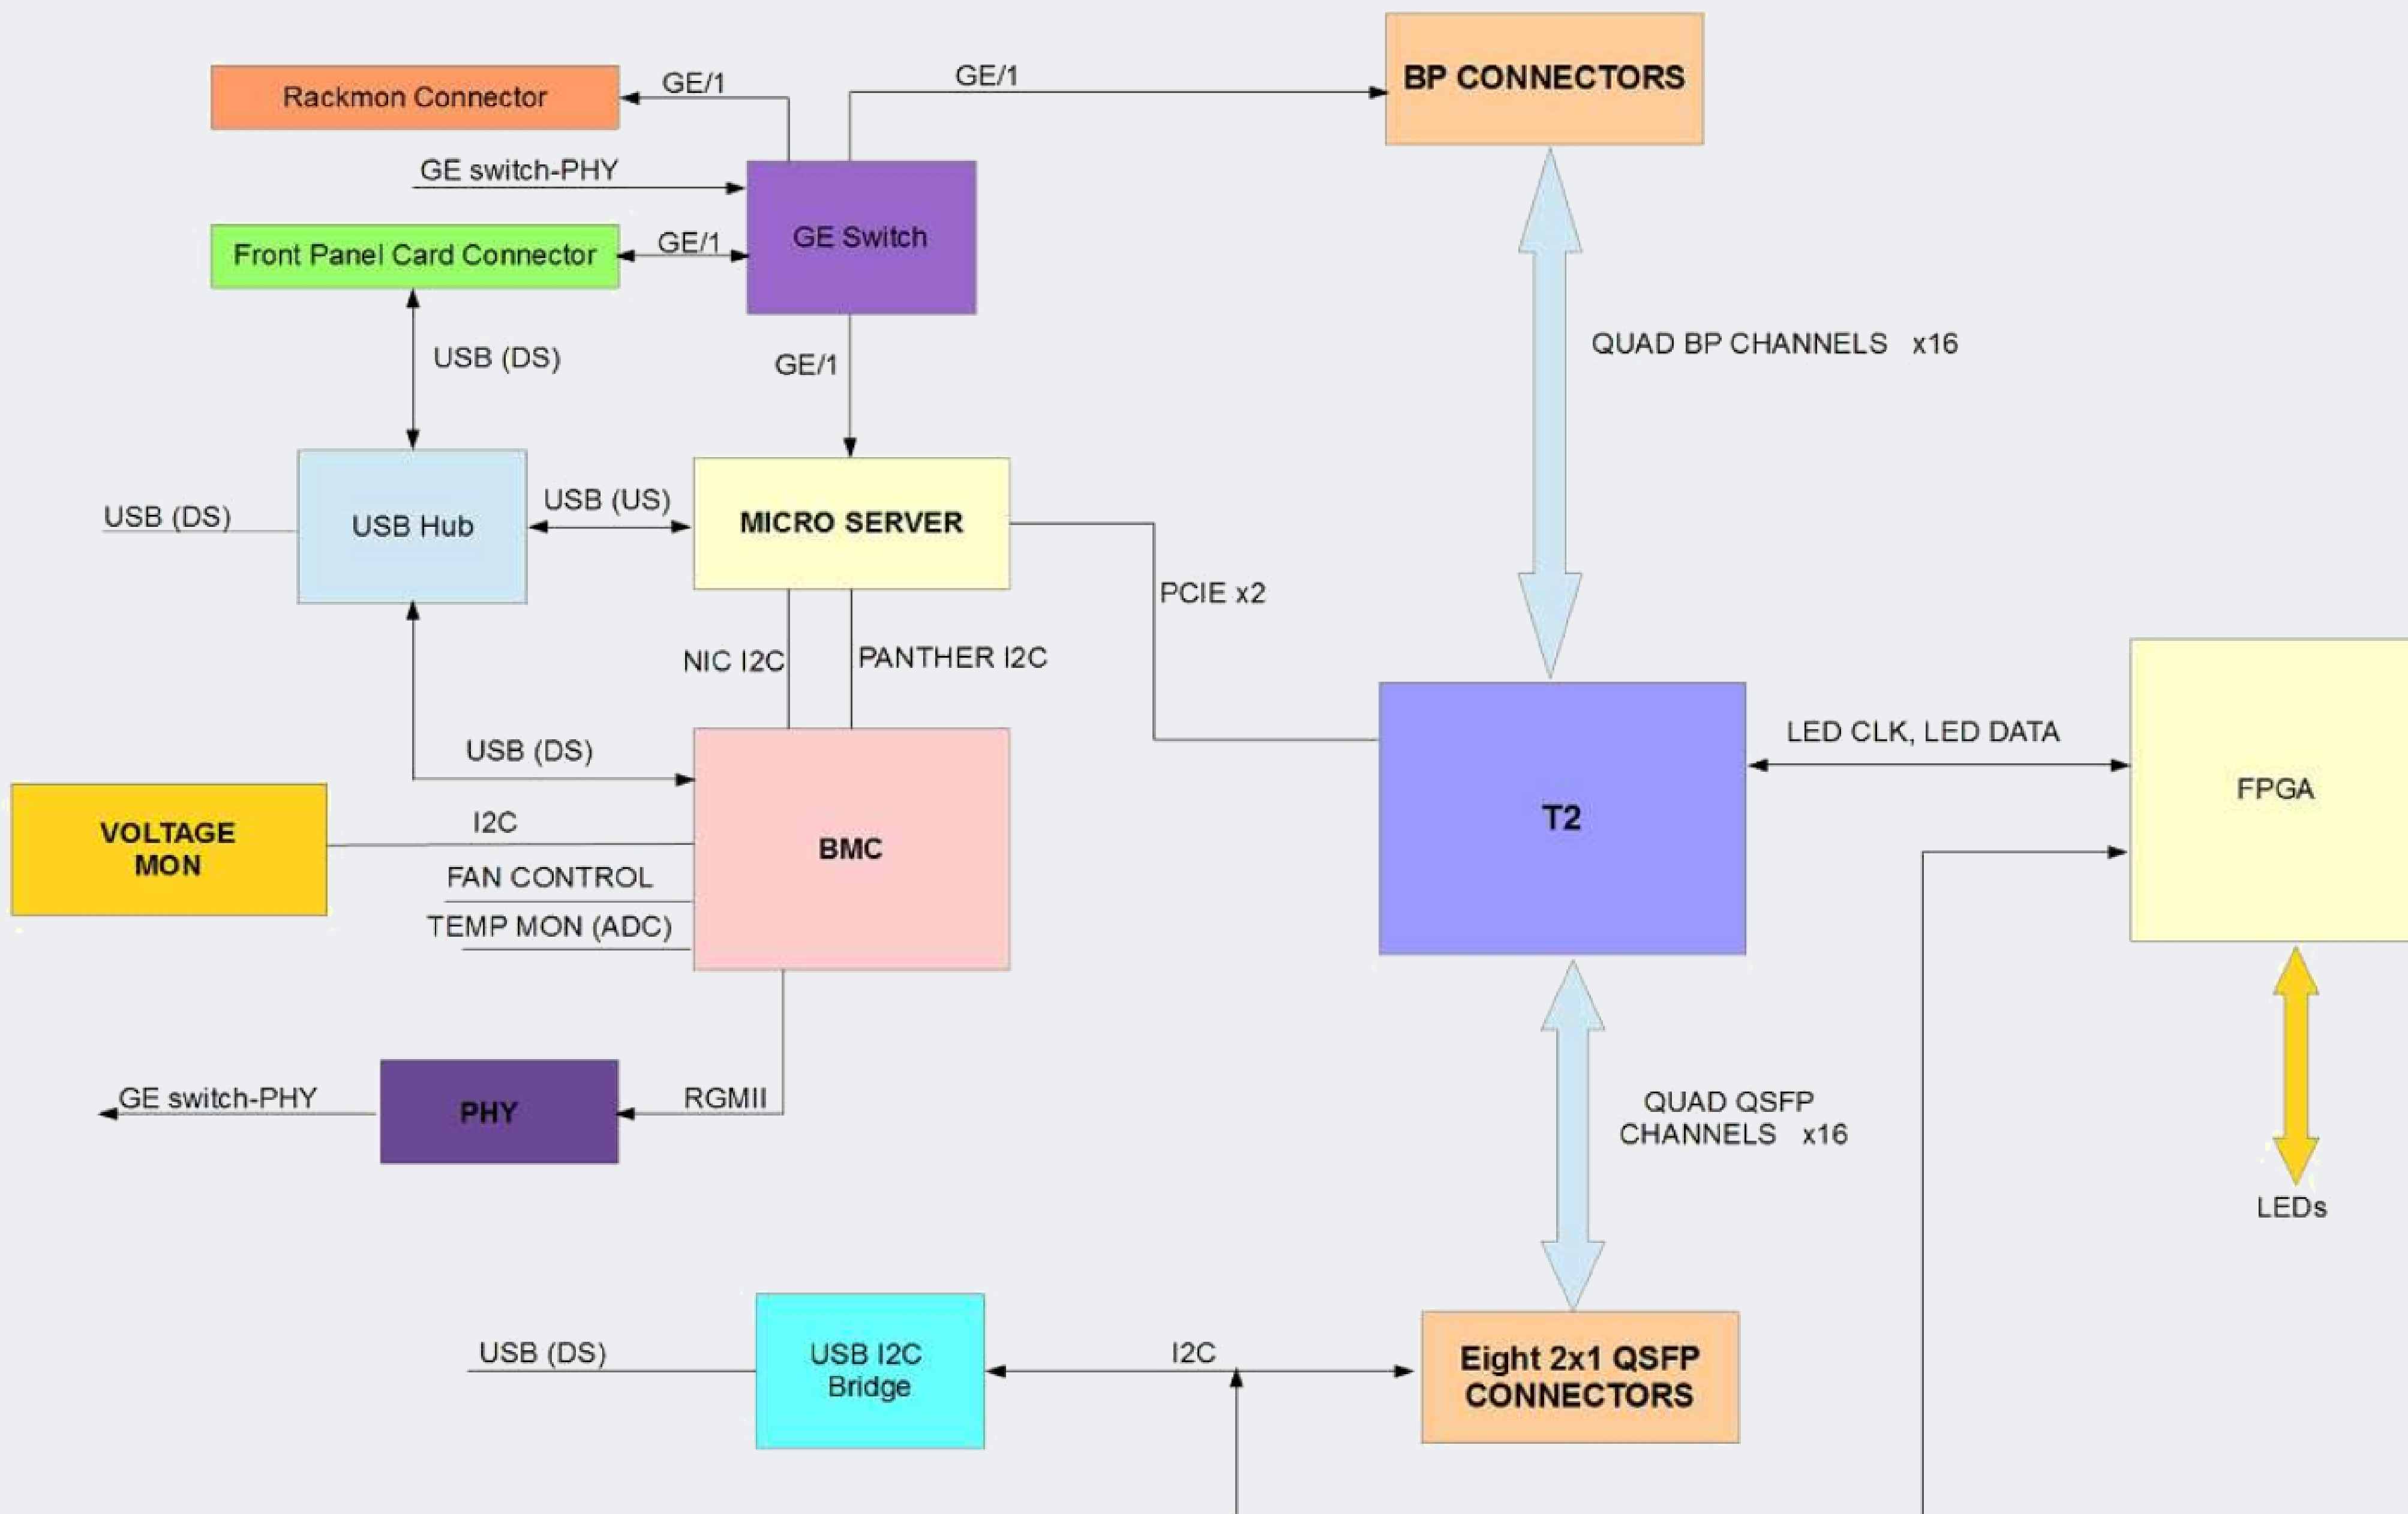
Feature Comparisons (Cont.)



	AST2500	AST2400	AST2300
Ethernet MAC	Dual MAC	Dual MAC	Dual MAC
Ethernet MAC Speed	10/100/1000M	10/100/1000M	10/100/1000M
USB 2.0 Controller (x1)	Yes (Hub)	Yes (Hub)	Yes (Hub)
USB 1.1 Controller (x1)	Yes (Hub)	Yes (Hub)	Yes (Hub)
USB Host Controller	Yes, x2 (2.0)	Yes, x1 (2.0)	Yes, x1 (1.1)
Hash & Crypto Engine	Hash/Crypto (DES/3DES/RSA)	Hash/Crypto (DES/3DES/RSA)	Hash/Crypto (DES/3DES/RSA)
VGA Resolution	1920x1200	1920x1200	1920x1200
KVM Resolution	1920x1200	1920x1200	1920x1200
UART	x5	x5	x5
Virtual UART	Yes	Yes	Yes
LPC Bus Controller	Master/Slave	Master/Slave	Master/Slave
GPIO Pins	232(max), 1.8V(x16)	216 (max)	152 (max)
Serial GPIO pins	80 (max)	80 (max)	64 (max)
MCTP over PCIe	Yes(with DMA)	Yes (with DMA)	Yes (w/o DMA)
Direct Port 80h to GPIOs	Yes (also SGPIO)	Yes	No





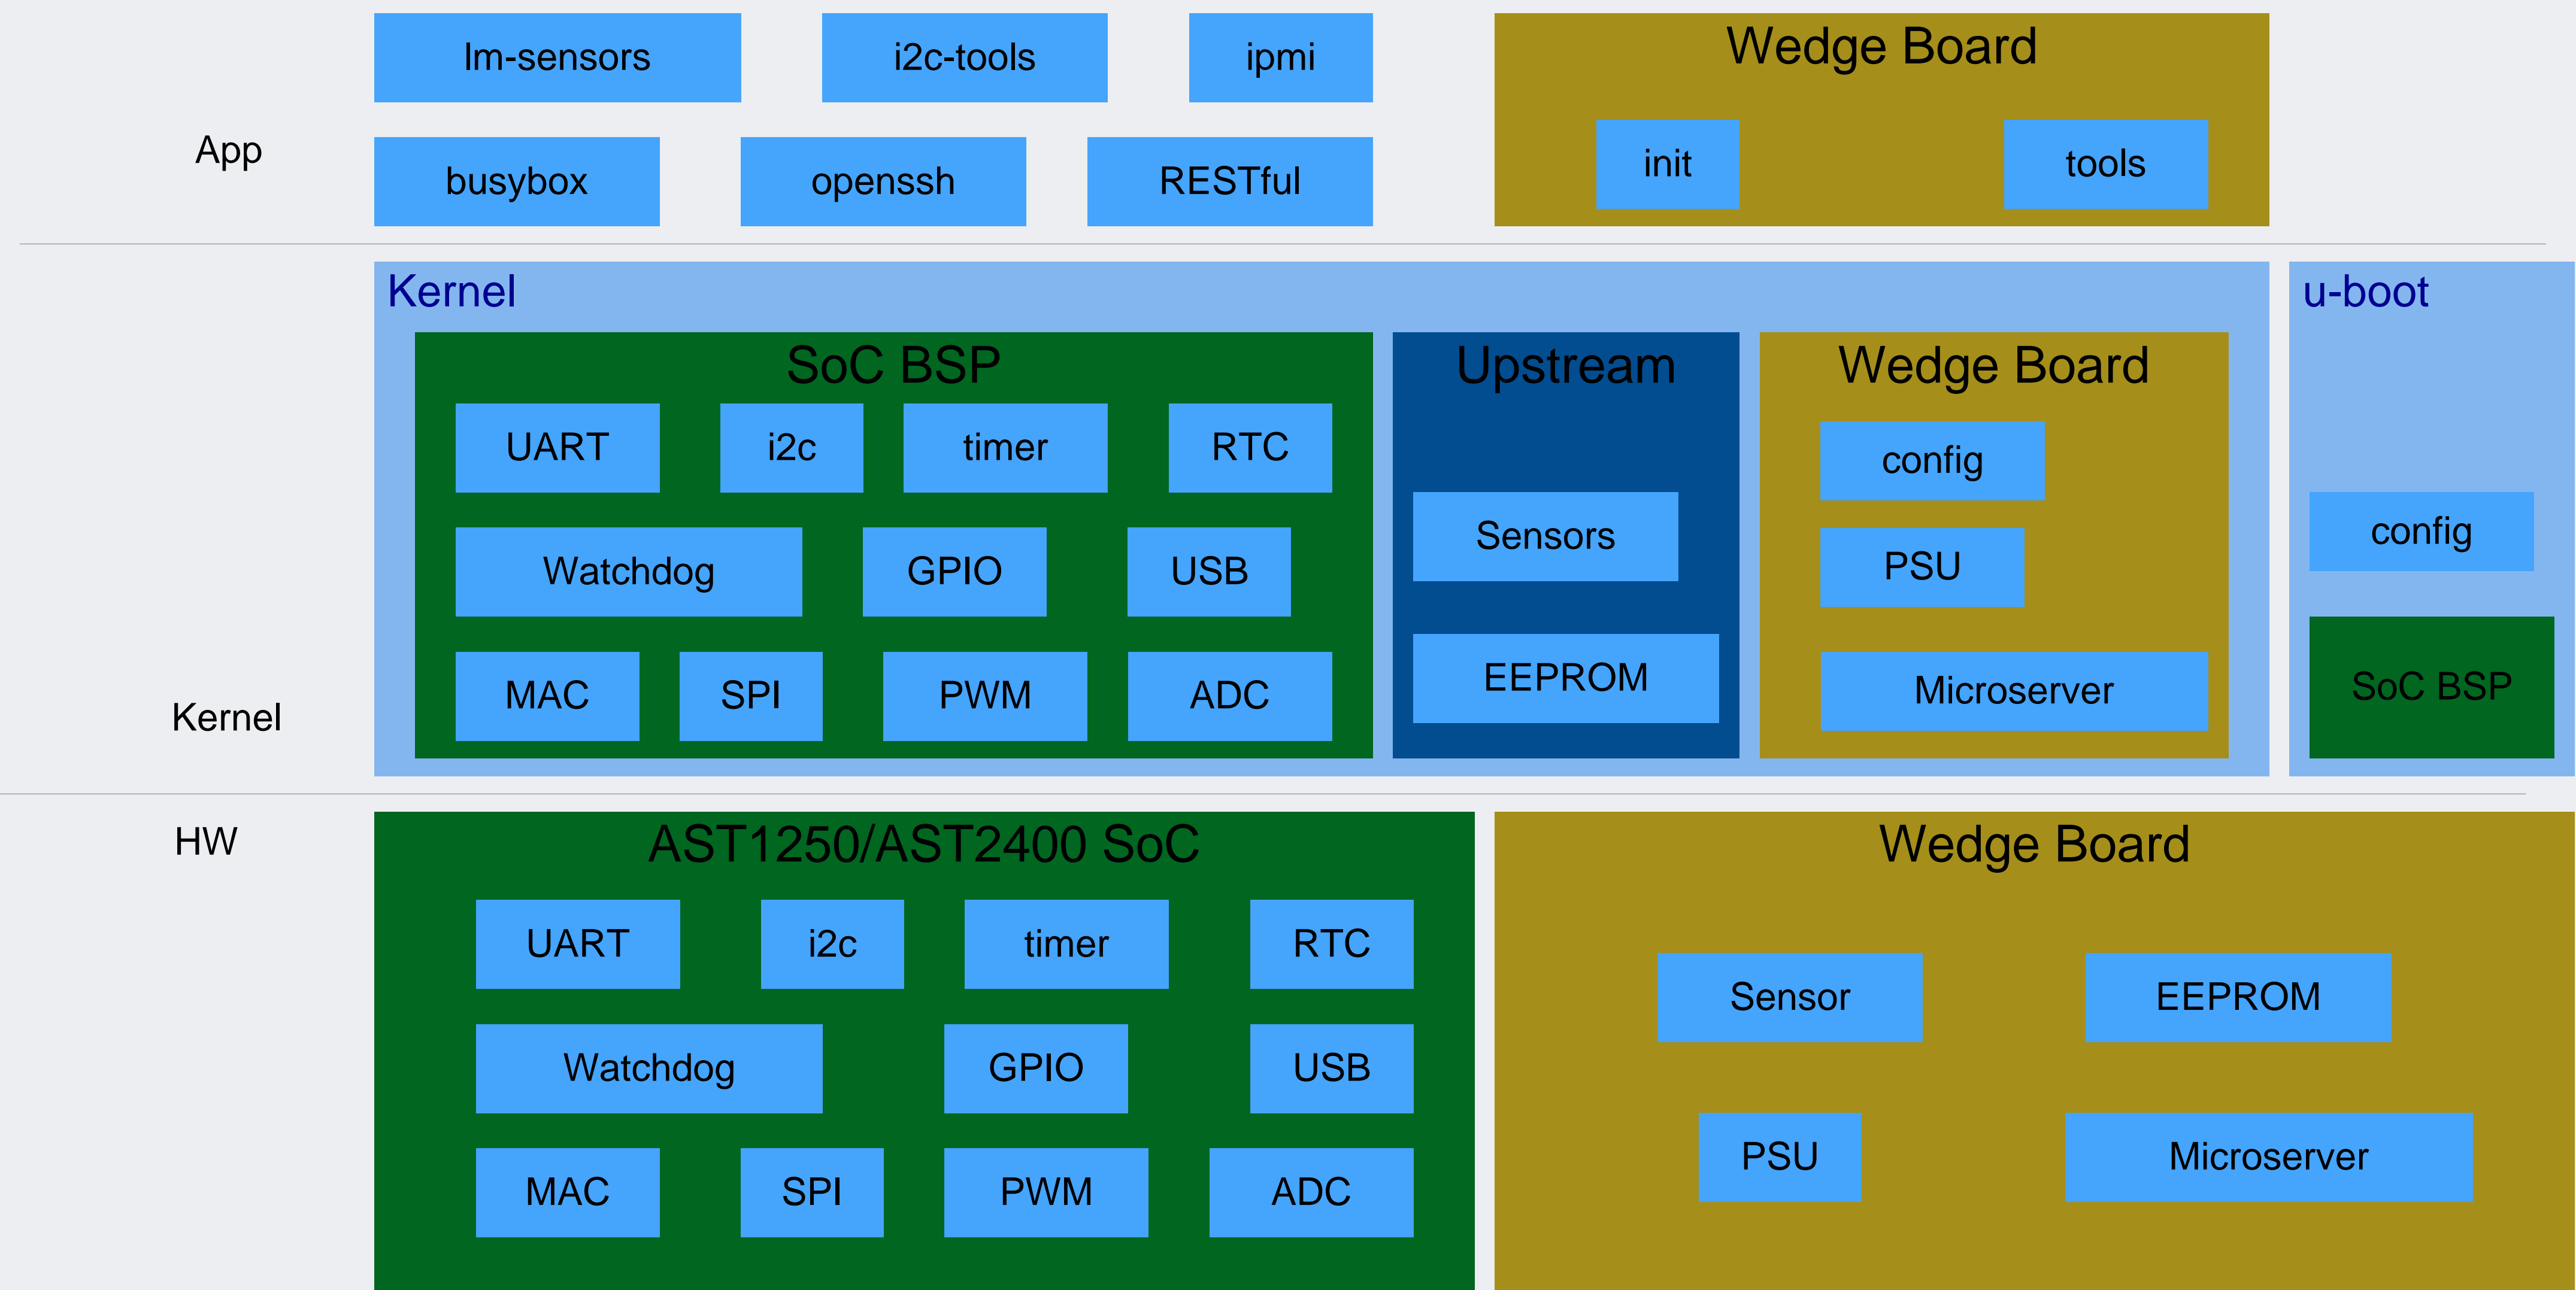


OpenBMC

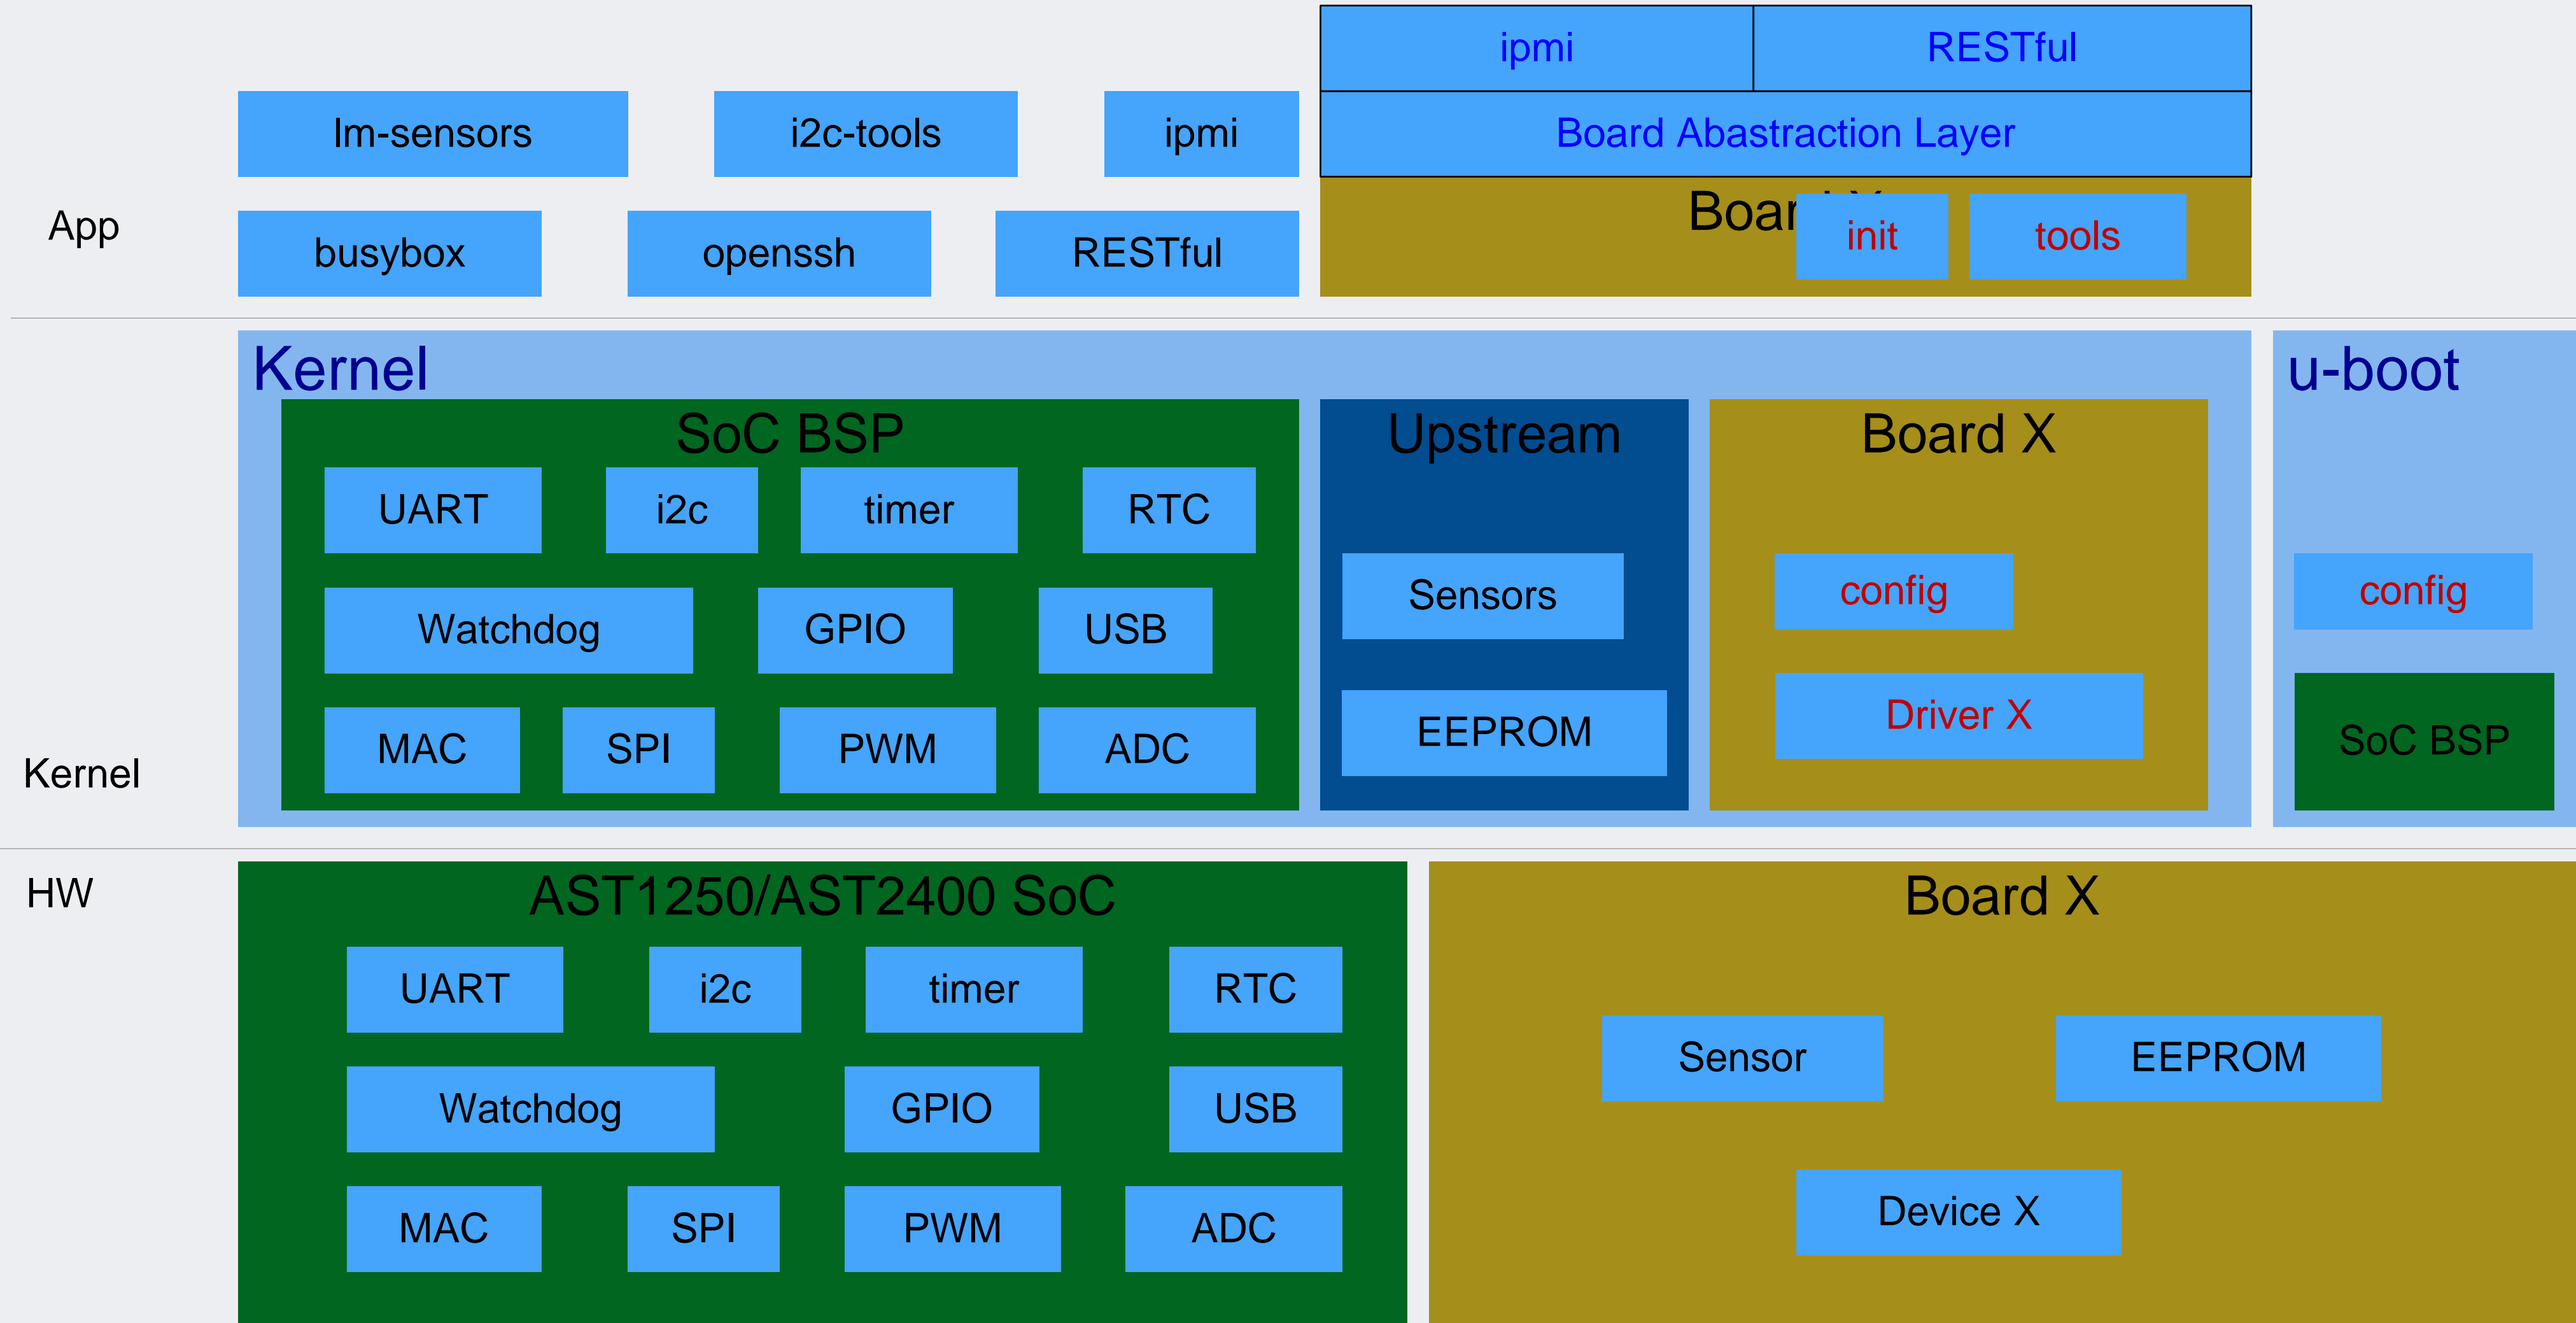
OpenBMC treats BMC as a regular server

OpenBMC is a customized Linux distribution

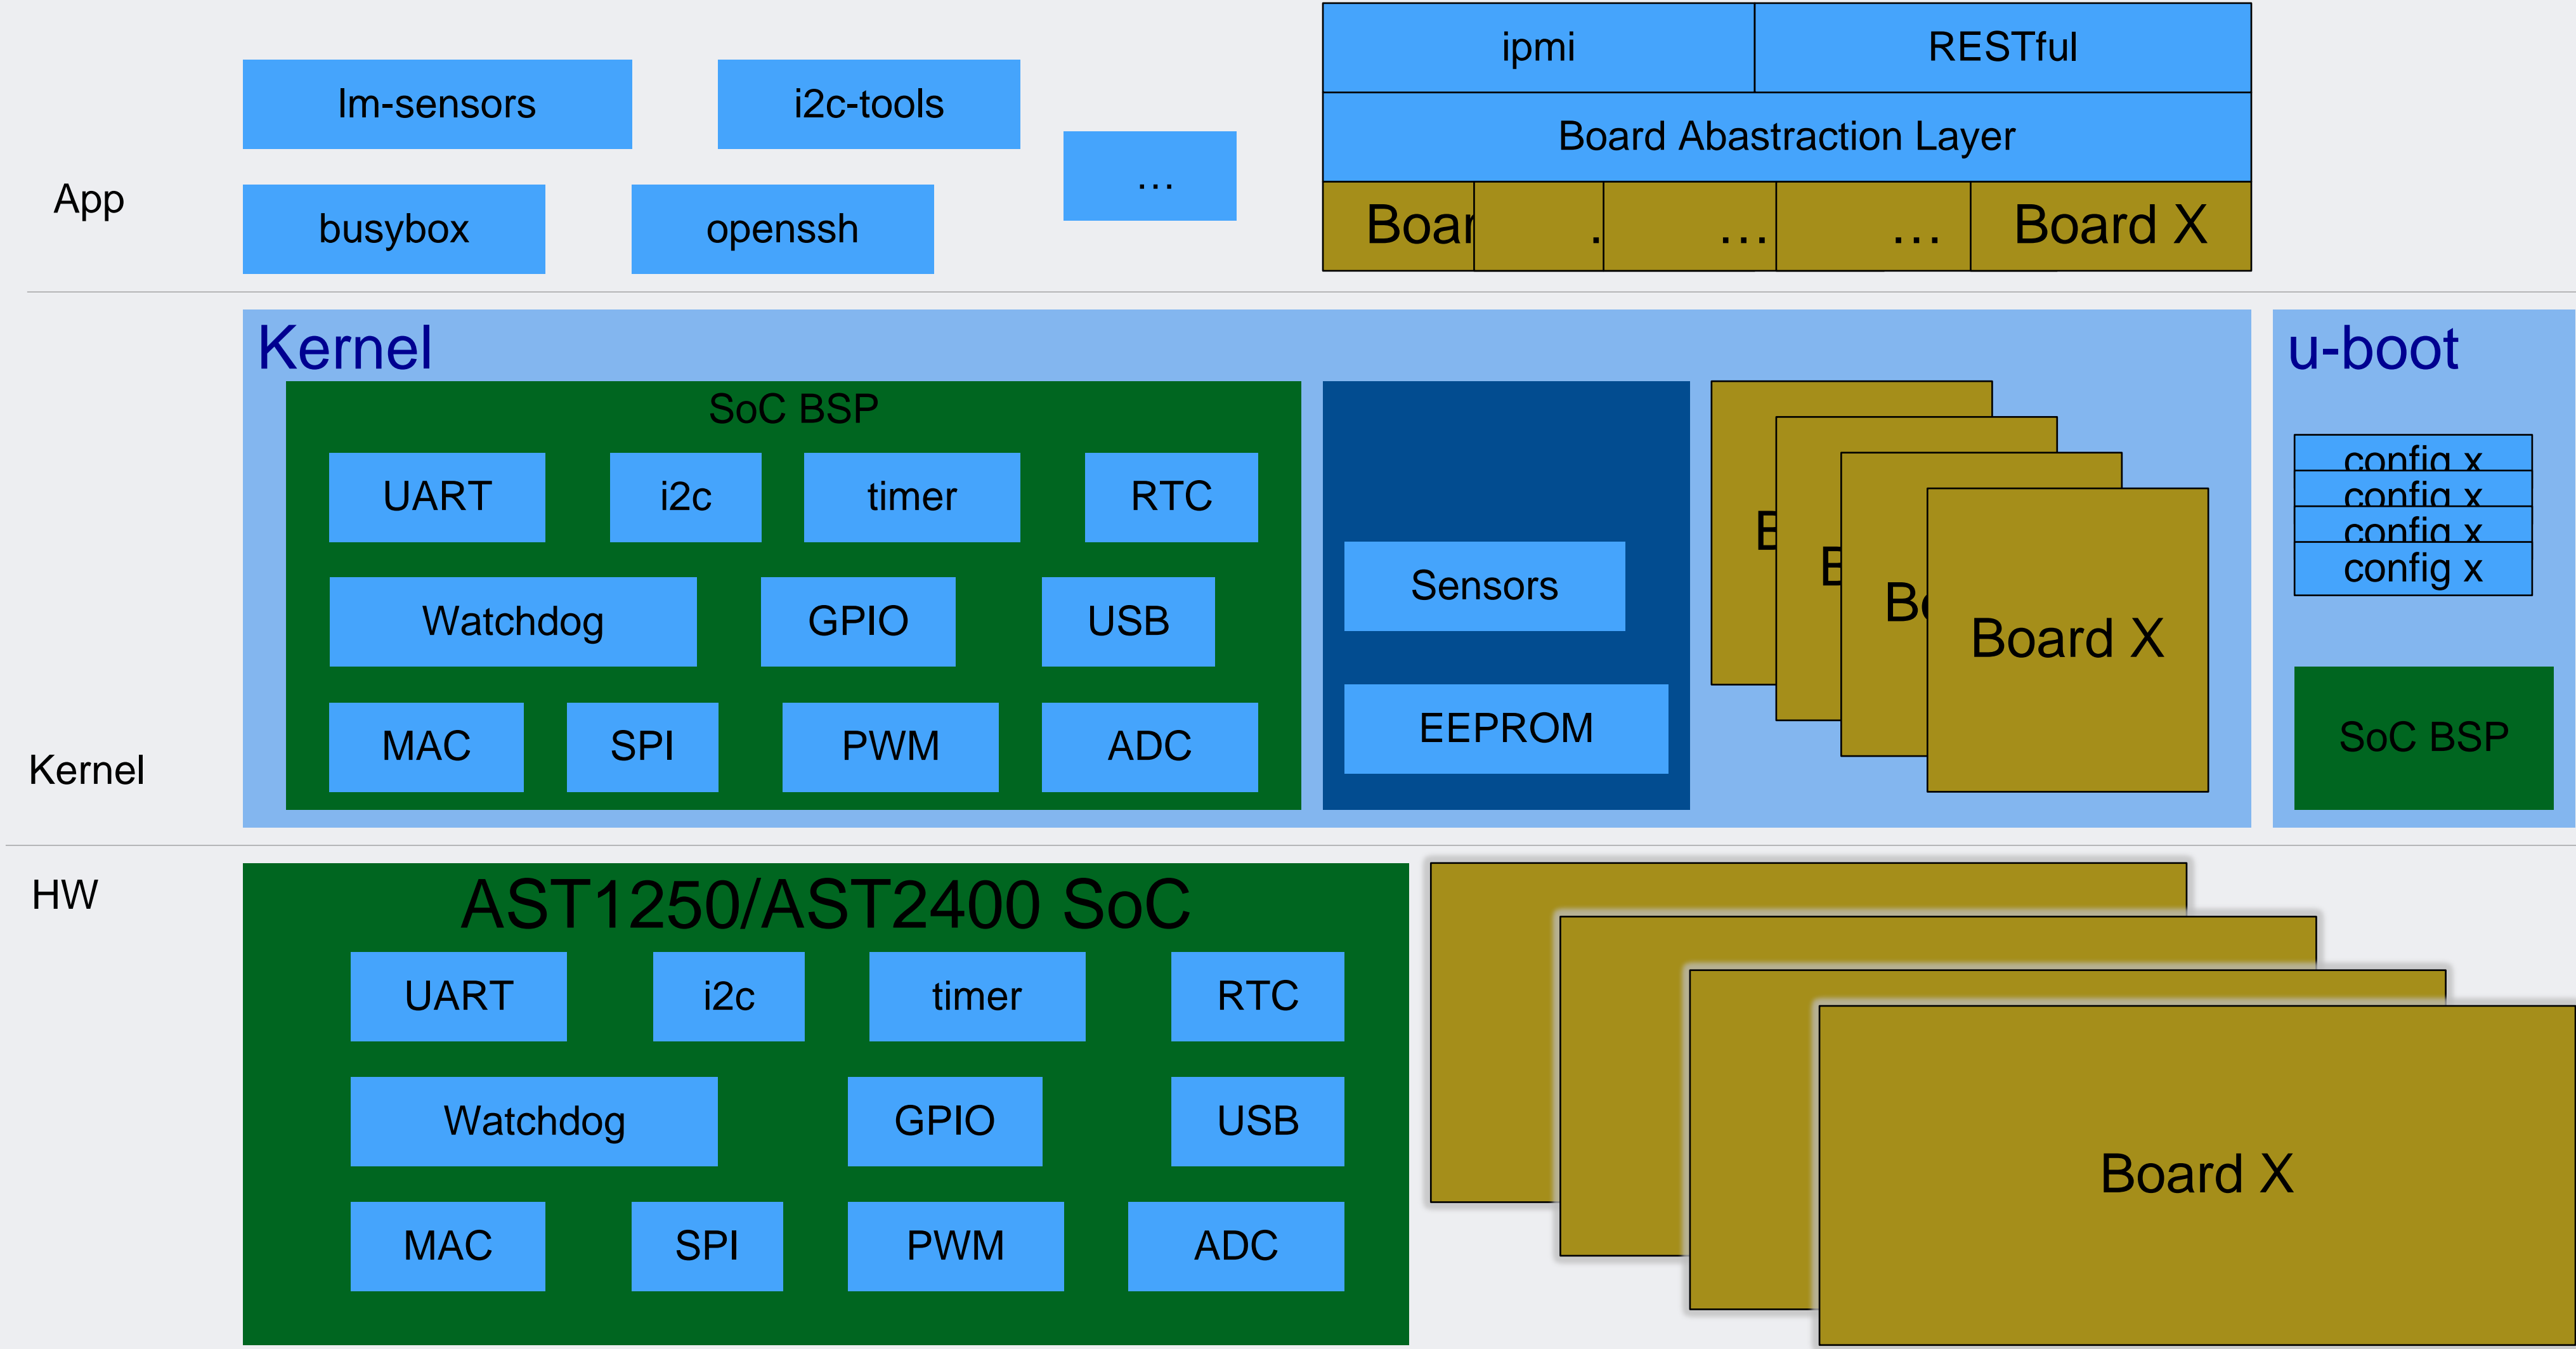
OpenBMC on Wedge



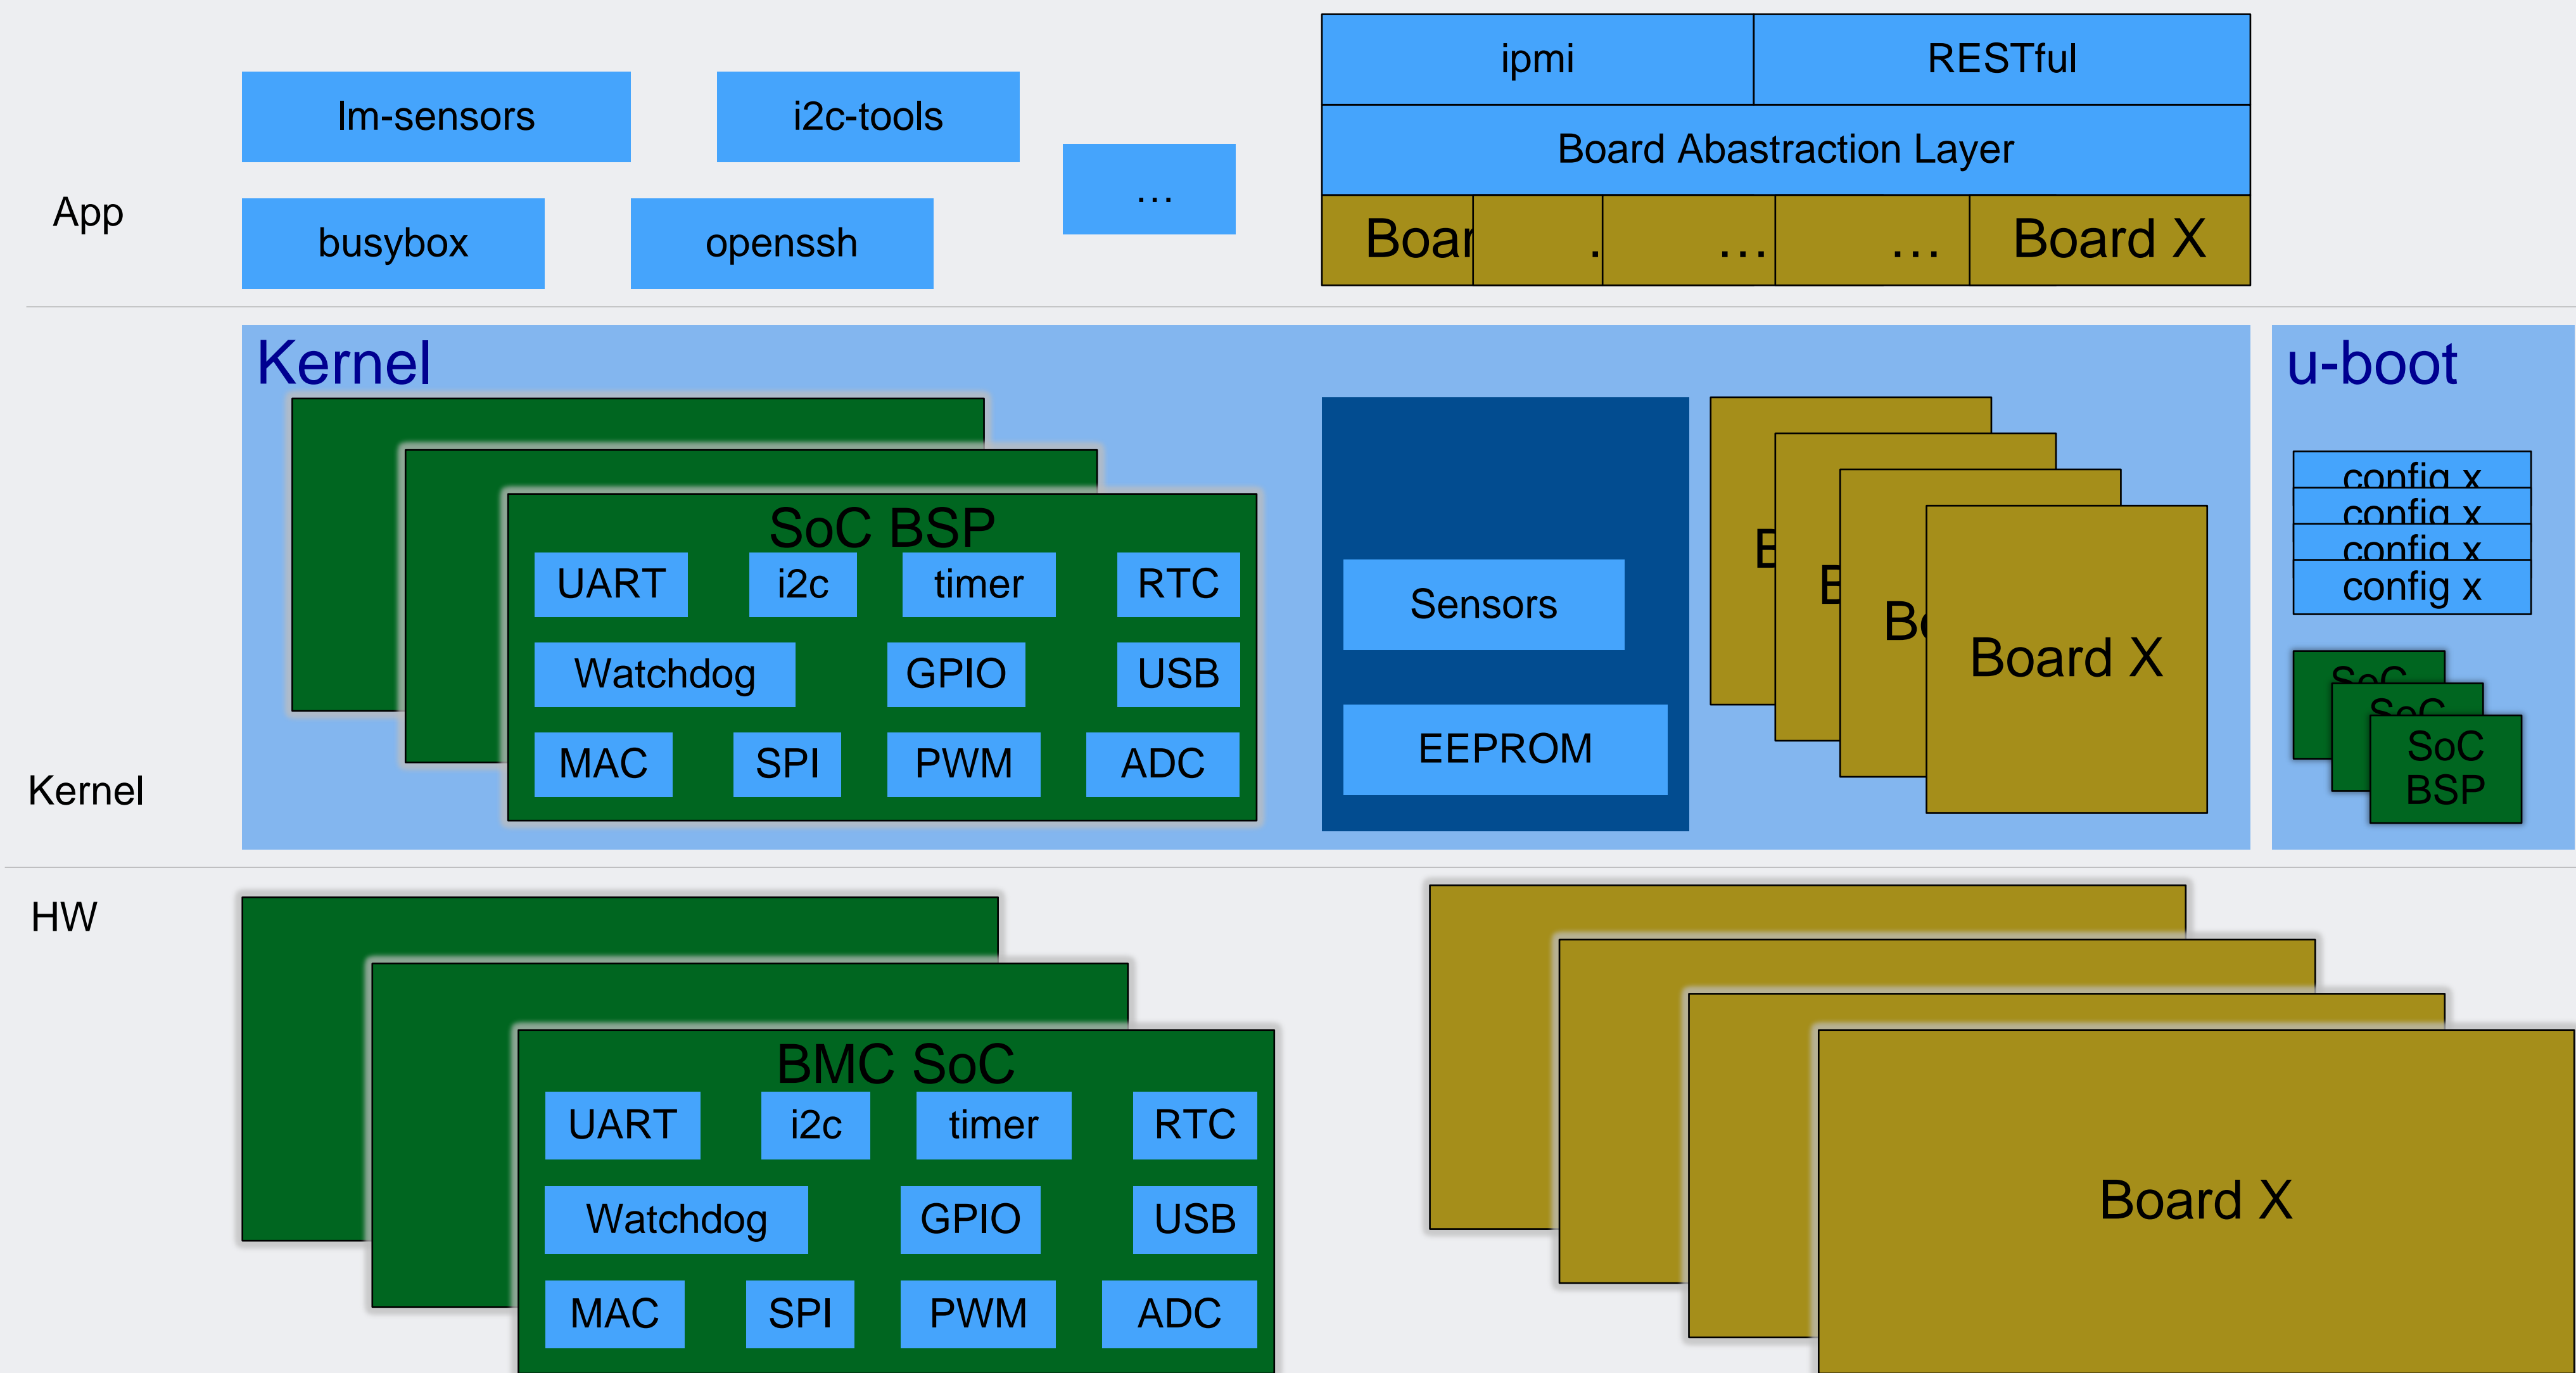
OpenBMC on board X



OpenBMC on different boards



OpenBMC



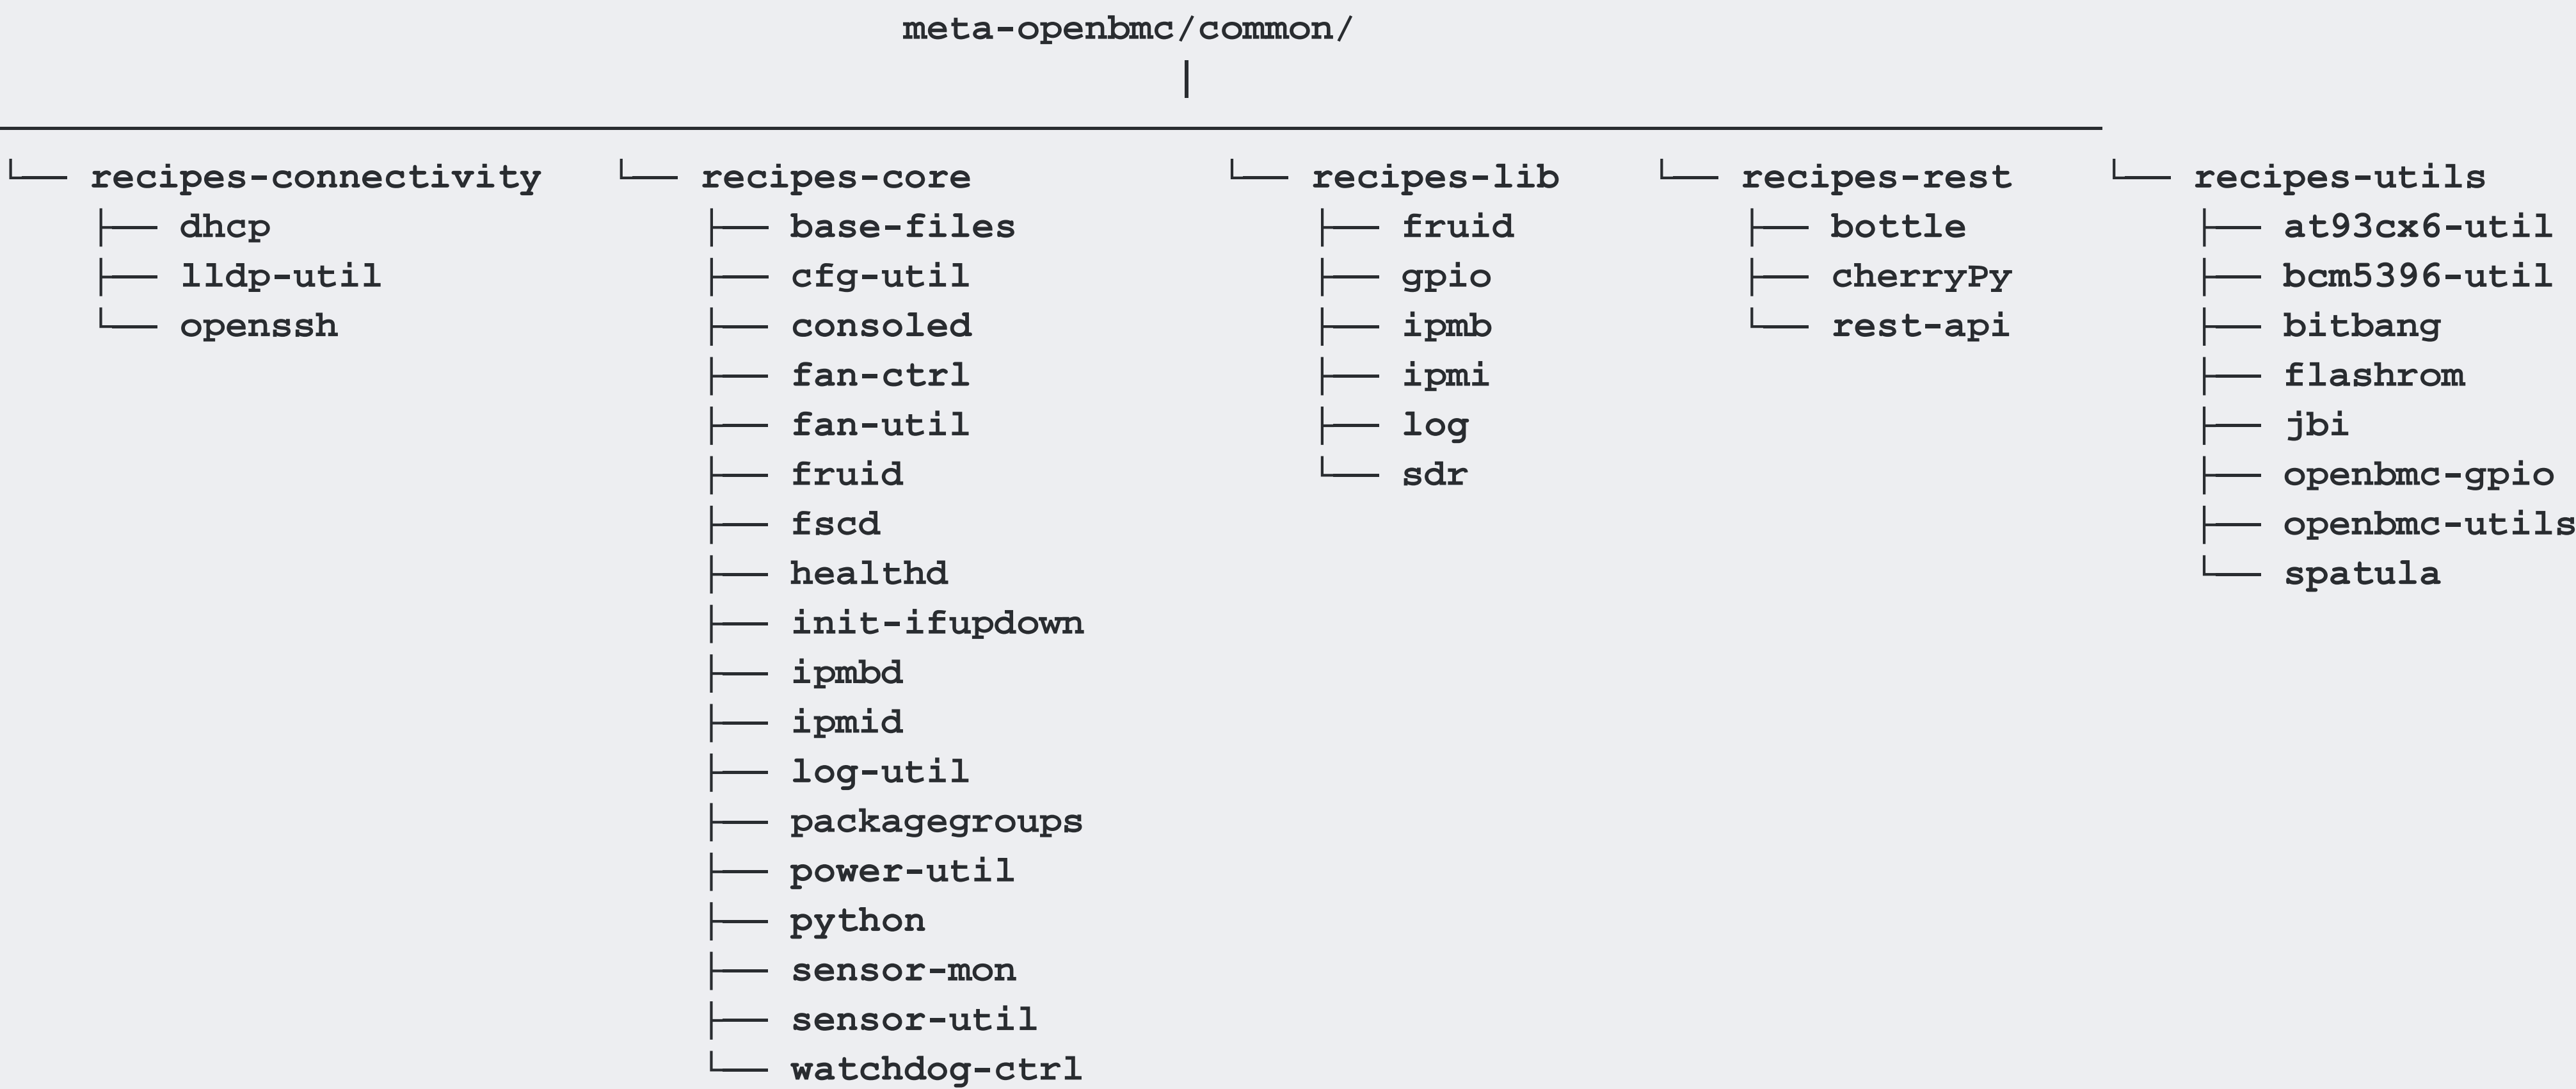
OpenBMC Source Code

Source code

- Use Yocto project
- 3 sets of layers
 - Common Layer
 - BMC SoC Layer
 - Board Specific Layer
- <https://github.com/facebook/openbmc>

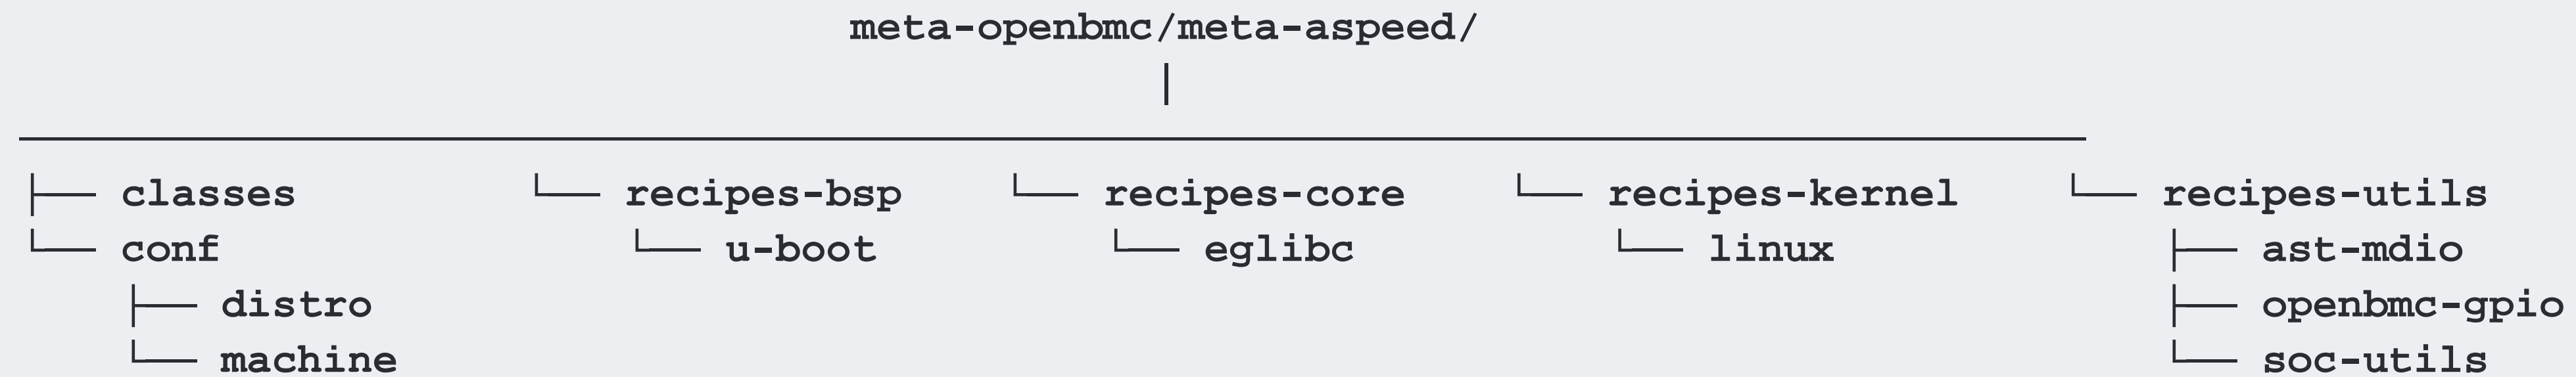
Common Layer

Common recipes and packages for all SoCs and boards



SoC Layer

SoC specific drivers and tols (bootloader, kernel, ...)



Board Specific Layer

Board specific customization, initialization, and tools



Layer architecture is empowered by Yocto

Code example

lm-sensors configuration

[meta-facebook/meta-yosemite/recipes-yosemite/lm_sensors/lmsensors %.bbappend:](#)

```
FILESEXTRAPATHS_prepend := "${THISDIR}/files:"
```

```
SRC_URI += "file:///yosemite.conf \  
            "
```

```
do_install_append() {  
    install -d ${D}${sysconfdir}/sensors.d  
    install -m 644 ../yosemite.conf ${D}${sysconfdir}/sensors.d/yosemite.conf  
}
```

[meta-facebook/meta-wedge/recipes-wedge/lm_sensors/lmsensors %.bbappend:](#)

```
FILESEXTRAPATHS_prepend := "${THISDIR}/files:"
```

```
SRC_URI += "file:///wedge.conf \  
            "
```

```
do_install_board_config() {  
    install -d ${D}${sysconfdir}/sensors.d  
    install -m 644 ../wedge.conf ${D}${sysconfdir}/sensors.d/wedge.conf  
}
```

```
do_install_append() {  
    do_install_board_config  
}
```


Code example

libpal

common/recipes-core/fruid/files/fruid-util.c:

```
...  
    ret = pal_get_fruid_eeprom_path(fru, eeprom_path);  
...
```

meta-facebook/meta-lightning/recipes-lightning/fblibs/files/pal/pal.c:

```
int  
pal_get_fruid_eeprom_path(uint8_t fru, char *path) {  
    return lightning_get_fruid_eeprom_path(fru, path);  
}
```

meta-facebook/meta-yosemite/recipes-yosemite/fblibs/files/pal/pal.c:

```
int  
pal_get_fruid_eeprom_path(uint8_t fru, char *path) {  
    return yosemite_get_fruid_eeprom_path(fru, path);  
}
```

Code example

openbmc-gpio

[meta-aspeed/recipes-utils/openbmc-gpio/files/ast2400_gpio_table.py:](#)

```
soc_gpio_table = {
    'A1': [
        Function('SD1WP#', BitsEqual(0x90, [0], 0x1)),
        Function('SDA13', BitsEqual(0x90, [26], 0x1)),
        Function('GPIOC7', None)
    ],
    'A10': [
        Function('GPIOU0', BitsEqual(0xa0, [8], 0x1)),
        Function('RMII2TXD0', BitsEqual(0x70, [7], 0x0)),
        Function('RGMII2TXD0', None)
    ],
    ...
}
```

[meta-facebook/meta-wedge100/recipes-utils/openbmc-gpio/files/board_gpio_table_v1.py:](#)

```
board_gpio_table_v1 = [
    BoardGPIO('GPIOB0', 'PANTHER_I2C_ALERT_N'),
    BoardGPIO('GPIOB1', 'MSERV_NIC_SMBUS_ALERT_N'),
    BoardGPIO('GPIOB2', 'DEBUG_PORT_UART_SEL_N'),
    ...
]
```

[common/recipes-utils/openbmc-gpio/files/openbmc_gpio_table.py:](#)

```
def setup_board_gpio(soc_gpio_table, board_gpio_table, validate=True):
```

Operation

Access

```
tfang@devbig453: ~$ ssh root@rswlmm-oob.13.atn1
```

```
Last login: Mon Mar 28 03:27:34 2016 from 2401:db00:11:d03a:face:0:1d:0
```

```
root@rswlmm-oob:~# dmesg|head -n 10
```

```
Linux version 2.6.28.9 (tfang@devbig288.prn2.facebook.com) (gcc version 4.8.2 (GCC) ) #1 Tue Sep 22  
17:05:20 PDT 2015
```

```
CPU: ARM926EJ-S [41069265] revision 5 (ARMv5TEJ), cr=00053177
```

```
CPU: VIVT data cache, VIVT instruction cache
```

```
Machine: AST2400
```

```
Memory policy: ECC disabled, Data cache writeback
```

```
On node 0 totalpages: 65536
```

```
free_area_init_node: node 0, pgdat c034e70c, node_mem_map c0373000
```

```
Normal zone: 512 pages used for memmap
```

```
Normal zone: 0 pages reserved
```

```
Normal zone: 65024 pages, LIFO batch:15
```

Upgrade

```
tfang@devbig453: ~$ scp kernel.v25.1 rootfs.v25.1 root@rswlmm-oob.13.atn1:/tmp
kernel.v25.1                                     100% 1685KB   1.7MB/s   1.7MB/s   00:01
rootfs.v25.1                                    100% 8994KB   1.5MB/s   2.4MB/s   00:06
tfang@devbig453: ~$ ssh root@rswlmm-oob.13.atn1
Last login: Mon Mar 28 03:35:34 2016 from 2401:db00:11:d03a:face:0:1d:0
root@rswlmm-oob:~# cd /tmp
root@rswlmm-oob:/tmp# cat /proc/mtd
dev:      size  erasesize  name
mtd0: 00060000 00010000 "u-boot"
mtd1: 00020000 00010000 "env"
mtd2: 00200000 00010000 "kernel"
mtd3: 00c00000 00010000 "rootfs"
mtd4: 01100000 00010000 "data0"
root@rswlmm-oob:/tmp# flashcp -v kernel.v25.1 /dev/mtd2 && flashcp -v rootfs.v25.1 /dev/mtd3 && reboot
Erasing block: 27/27 (100%)
Writing kb: 1684/1684 (100%)
Verifying kb: 1684/1684 (100%)
Erasing block: 141/141 (100%)
Writing kb: 8993/8993 (100%)
Verifying kb: 8993/8993 (100%)

The system is going down for reboot NOW!3.atn1.facebook.com (pts/0) (Mon Mar
root@rswlmm-oob:/tmp# Connection to rswlmm-oob.13.atn1 closed by remote host.
Connection to rswlmm-oob.13.atn1 closed.
tfang@devbig453: ~$
```

Sensors

```
root@rswlmm-oob:~# sensors |grep Temp
```

```
CPU Temp:      +15.5 C
```

```
DIMM0 Temp:    +29.5 C
```

```
Inlet Temp:    +22.5 C  (high = +80.0 C, hyst = +75.0 C)
```

```
Switch Temp:  +26.0 C  (high = +80.0 C, hyst = +75.0 C)
```

```
Outlet Temp:  +27.5 C  (high = +80.0 C, hyst = +75.0 C)
```

```
root@rswlmm-oob:~# sensors |grep Fan
```

```
Fan 2 front: 7114 RPM
```

```
Fan 3 front: 7020 RPM
```

```
Fan 1 front: 6939 RPM
```

```
Fan 0 front: 7118 RPM
```

```
Fan 2 rear: 4611 RPM
```

```
Fan 3 rear: 4750 RPM
```

```
Fan 1 rear: 4760 RPM
```

```
Fan 0 rear: 4389 RPM
```


RESTful



```
tfang@devbig453: ~$ curl -s http://rswlmm-oob.13.atn1:8080/api/sys/mb/fruid| json_reformat
```

```
{
  "Information": {
    "Location on Fabric": "DGE",
    "Product Sub-Version": 0,
    "Facebook PCB Part Number": "",
    "CRC8": "0x0",
    "System Assembly Part Number": "135-000010-01",
    "Product Serial Number": "AF03041729",
    "System Manufacturing Date": "01-21-15",
    "Local MAC": "70:72:CF:F1:9B:A4",
    "Assembled At": "Joytech",
    "ODM PCBA Serial Number": "AF03041729",
    "Product Asset Tag": "2274855",
    "Product Name": "Wedge-DC-F",
    "ODM PCBA Part Number": "CP2FB6632003A",
    "Product Production State": 3,
    "Product Part Number": "20-001102",
    "PCB Manufacturer": "Hitachi",
    "System Manufacturer": "Joytech",
    "Extended MAC Address Size": 17751,
    "Facebook PCBA Part Number": "132-000010-01",
    "Version": 0,
    "Extended MAC Base": "70:72:CF:F1:9B:A5",
    "Product Version": 1
  },
  "Resources": [
    ],
  "Actions": [
    ]
}
```

Status

Deployment









Next

Next steps

- Challenges
 - Driver stability
 - Better tooling: development, provisioning, monitoring
 - HW availability
- OpenBMC Developer Platform
 - Baseboard with AST2500 BMC
 - COM Express connectivity: Choose and plugin host CPU
 - Portwell COMe OpenBMC carrier board

Q/A

facebook